



Simplified, Parallelized InSAR Scientific Computing Environment

Paul A. Rosen, Eric Gurrola, Piyush Agram,
Geoffrey Gunter
Jet Propulsion Laboratory
California Institute of Technology

Michael Aivazis
Parasim, Inc.


Howard Zebker
Stanford University

Earth Science Technology Forum
NASA Ames Conference Center
June 12, 2019




Simplified Parallelized InSAR Scientific Computing Environment (SPISCE) Background

- Past ESTO investments have created the InSAR Scientific Computing Environment (ISCE), used by individual users and semi-operational production systems (ARIA)
- The growth in data from Sentinel 1 and soon NISAR requires a fundamentally different way of interacting with data
- This work exploits the power for GPUs to slice through data and implement framework elements to give scientists a different user experience



GPU-accelerated Geocoded SLC data and related workflows will help users create derived products with simple editable operations



Objectives

- Develop a workflow that generates single look complex (SLC) images in an easily ingestible form for data utility, incorporating all necessary InSAR phase corrections and resampling to provide products that anyone can understand
- Develop methods of computational acceleration by exploiting back projection methods on cloud-enabled GPU platforms to directly compute focused and InSAR-corrected imagery in UTM (Landsat grid).
- GPU-parallelize and otherwise accelerate key algorithms in InSAR Scientific Computing Environment (ISCE) to gain factors of 10-100 increase in speed.
- Develop framework technologies at the user interface that support a more natural way for scientists to specify products and actions.
- Provide open source solutions



Summary of Results

Geocoded SLC workflow

- Developed the workflow – foundation of data processing in Stanford's research group
- InSAR process flow experimentation
 - Is it needed to create redundant interferograms to get best results in the context of the starting point of geocoded SLC?
- Socializing geocoded SLC's with the community
 - Papers at AGU
 - Papers at LPS and IGARSS this year.

GPU parallelization of back-projection processor

- Significant (10-100 X) improvements to speed and efficiency – 2 seconds per image

GPU parallelization of ampcor processor



- Excellent (1000 X) improvement in speed and efficiency

Pyreized workflow

- ISCE 2 components upgrade to pyre – “complete”
- SPISCE 1.0 with GUI showing each component as a workflow node near complete

User-friendly products

- Creating a time series view of Kilauea eruption is greatly simplified by starting with geocoded SLCs facilitated by backprojection

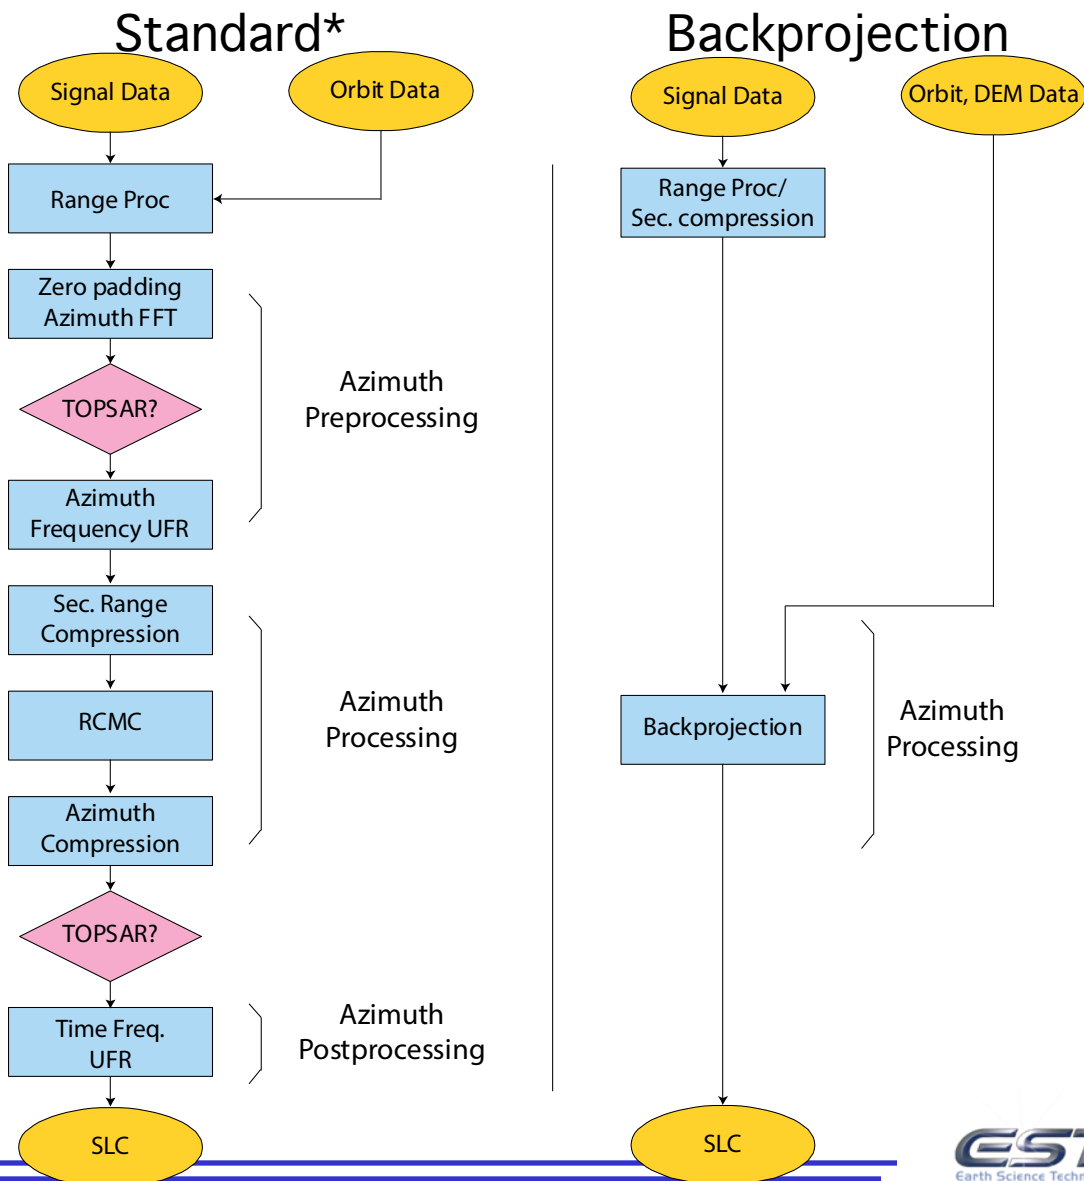


Traditional
“range-doppler”
Image – hard to
interpret!



L0 SLC processing for TOPS

- Dealiasing computationally expensive and very intricate though elegant
- Backprojection same for all modes
- Short integration time makes backprojection efficient

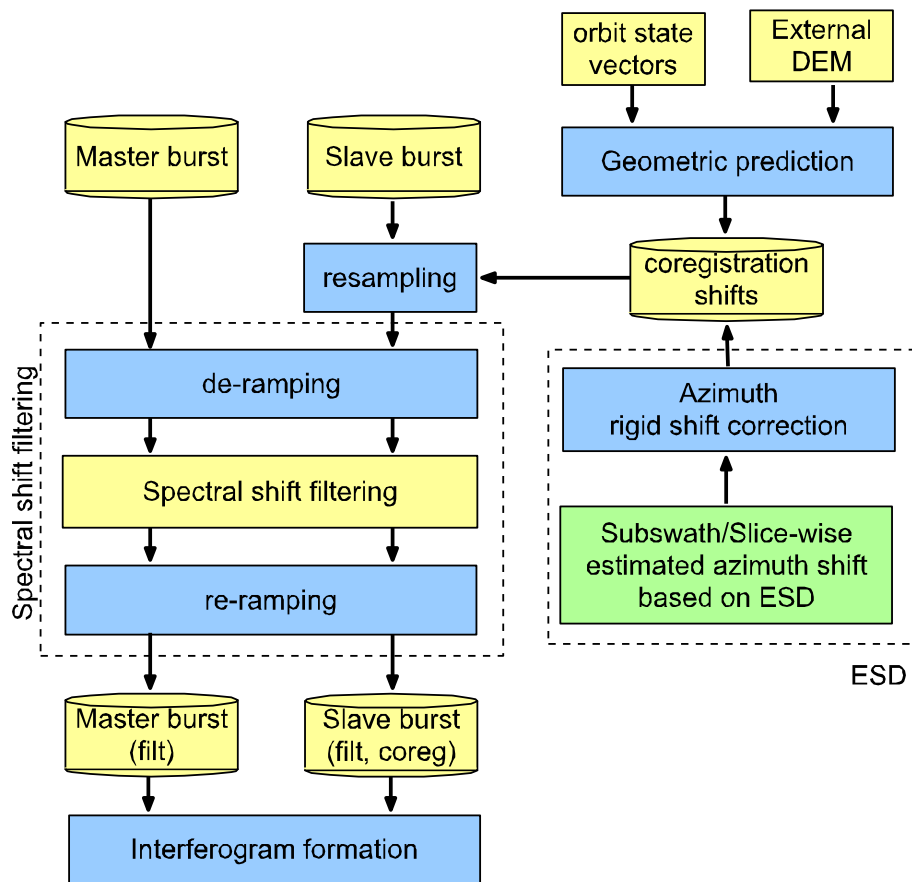


*Sentinel 1 SAR Technical Guide




Interferogram formation – by user

Standard (Yague-Martinez et al., 2016)



Backprojection






Example: Noise decreases, accuracy increases with “redundant” data generation


$$\phi_{ij} = \phi_{ij}^{defo} + \phi_{ij}^{aps} + \phi_{ij}^{decor}$$

Independent stacking



$$\phi_{ind} = \frac{1}{2}(\phi_{13} + \phi_{24})$$


Redundant stacking




$$\phi_{red} = \frac{1}{4}(\phi_{13} + \phi_{24} + \phi_{14} + \phi_{23})$$

Slow slip
earthquake
signal in
Cascadia

Independent stacking



Redundant stacking



Can now generate “redundant”
interferograms where before it was not
practical

Conclusion: *AIST improved workflow
makes InSAR sensitive and accurate for
many more events and studies*




Example of Socialization of Products with Community

Observed Deformation and Seismicity Application

From 1407 interferograms enabled by AIST backprojection processor


Induced seismicity from oil and gas production – Delaware Basin, Texas

Participants: Texas Bureau of Economic Geology, Stanford Center for Induced and Triggered Seismicity



At left: Seismicity (eq's) correlates well with observed deformation signal

At right: Production correlates well with observed deformation signal south of river



Neither seismicity nor deformation correlate spatially north of river, likely due to geological setting and fault type

<---- subsidence uplift ---->
-10 -8 -6 -4 -2 0 2 4 6 8 10
Weighted LOS Displacement (cm)


40
km

70 km




Image cross-correlation GPU-acceleration

- Slide a tile within the search window whose shape is the same as the reference tile
- Compute a correlation measure at each possible placement
- Record the placements of maximum correlation for each tile
- Build and post-condition a correlation surface to form the offset field



reference
image



target
image

C++ intermediary code improvements (20X)


- memory mapped IO
- algorithms from the C++ STL
- a dedicated data structure to simplify multi-dimensional indices arithmetic

GPU code improvements (addl. 1500X)



- careful packing of the input tiles in memory
- minimizing the data transfer cost between host and device
- eliminating memory bank conflicts within the device
- highly tuned implementation of the reduction kernels



GPU-Accelerated back projection comparison



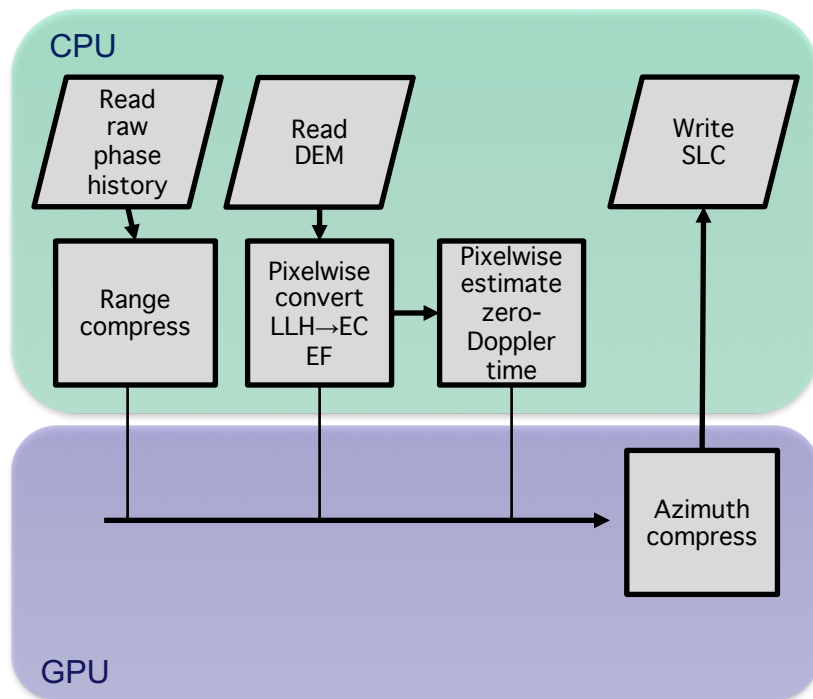
CPU: Intel Xeon CPU E5-2698 v4 | GPU: Nvidia
Tesla V100



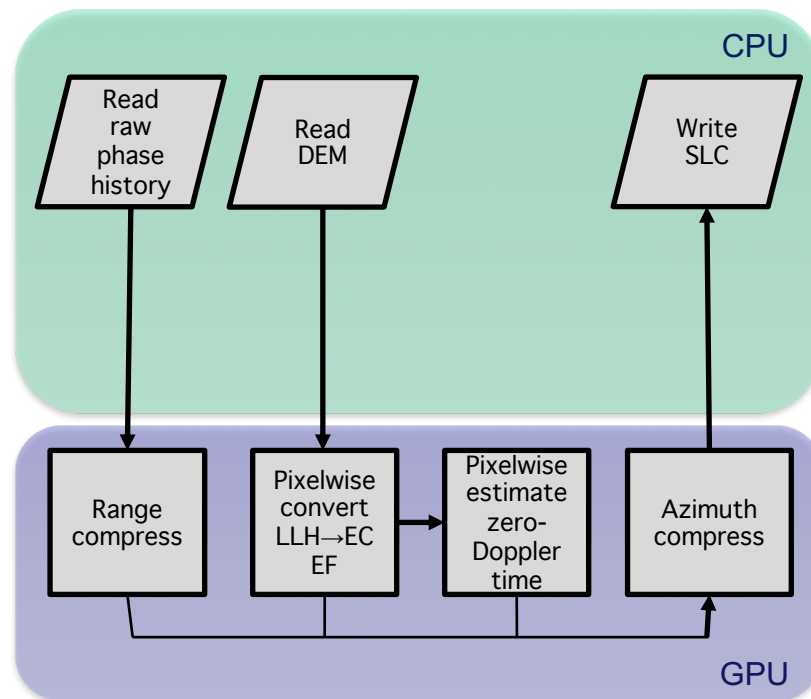


GPU backprojection code modifications

Old code



Modified code





Workflow visualizer (Flo)

- We are building a web-based application
 - simplify and organize access to the dozens of configurable parameters
 - support interactive processing
 - embedded within the isce application
- Managing the user experience for such a complex system is a challenge
 - multiple layers of information
 - complicated workflows tend to become difficult to render in an understandable way
 - users have high expectations from such interfaces
- Dozens of attempts to solve this problem
 - we think we have a good chance to avoid the known pitfalls, and provide a workable solution



Visual language prototype



atomic data



composite data



composite data with detail



nameless factory



named factory



n x m factory



raw input



final product





data bindings





Simple flow examples



Server:

- isce+pyre (GitHub: [isce-framework/isce2](https://github.com/isce-framework/isce2) and [pyre/pyre](https://github.com/pyre/pyre))

Client:

- SVG (and soon webgl) for the rendering of the graph
- ReactJS framework to build the user experience <https://reactjs.org>.





Client-Server interaction: a collection of data queries, mutations and subscriptions.


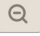

- GraphQL middleware (<https://graphql.org>).
- WebSockets.



Early design – editing factory metadata

Example Workflow NameSaved







Factories


Products


Q Search


▼


 Factory Name
Type · Topology · 2 Inputs · 3 Out...


 Factory Name
Type · Topology · 8 Inputs · 2 Out...


 Factory Name
Type · Topology · 5 Inputs · 1 Out...


 Factory Name
2 Inputs · 3 Outputs


 Factory Name
2 Inputs · 3 Outputs


 Factory Name
2 Inputs · 3 Outputs


 Factory Name
2 Inputs · 3 Outputs


 Factory Name
2 Inputs · 3 Outputs


 Factory Name
2 Inputs · 3 Outputs

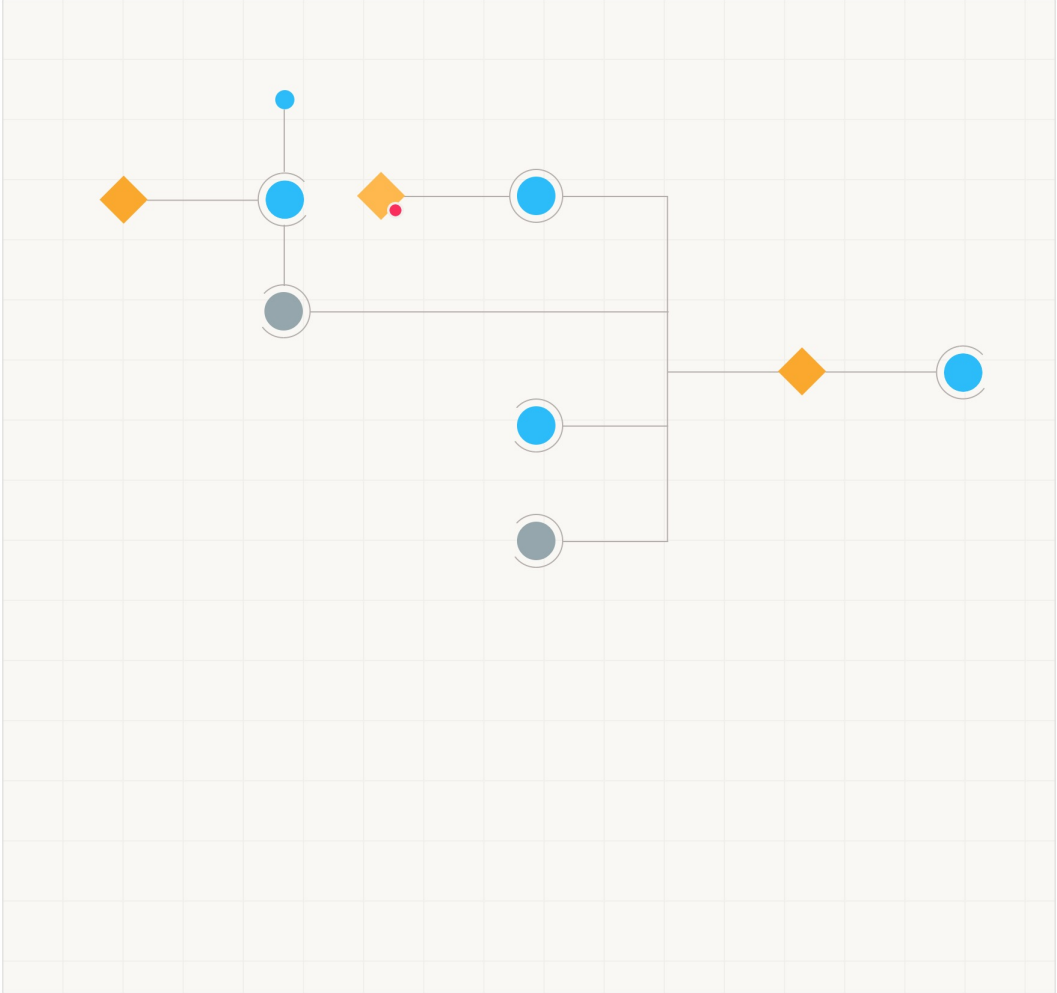
 Factory Name
2 Inputs · 3 Outputs


 Factory Name
2 Inputs · 3 Outputs

 Factory Name
2 Inputs · 3 Outputs

 Factory Name
2 Inputs · 3 Outputs

 Factory Name
2 Inputs · 3 Outputs



 isce.topography.dem.srtm3
2 Inputs · 3 Outputs

Label

Value

This is an error message.

Label

Placeholder

Label

Value

▼

☐ Label


Label

☐ Label

☐ Label

Integer

1



Float

1.0

Numeric Range - Integer

1

500

● Ready

Shortcuts · View Logs



Early design – flow execution

Example Workflow NameSaved

Factories

Products

Q Search

Y

Factory Name

Type · Topology · 2 Inputs · 3 Out...

Factory Name

Type · Topology · 8 Inputs · 2 Out...

Factory Name

Type · Topology · 5 Inputs · 1 Out...

Factory Name

2 Inputs · 3 Outputs

Factory Name

2 Inputs · 3 Outputs

Factory Name

2 Inputs · 3 Outputs

isce.topography.dem.srtm3

2 Inputs · 3 Outputs

Label

Placeholder

Label

Value

☐ Label

Label

isce.topography.dem.srtm3.n31w118

DEM Elevation Tile at n31 w118 · String · M...

2019-01-11T19:06:33.540400+00:00 heroku[router]: at=info method=GET path="/assets/logo-color-cb32dd3ef4021fd491c823578fc51a5ba1640a9cb1d441e3fc02d2939f9215d7.png" host=www.crowdvest.co request_id=ed64f91c-3c9d-4c2c-ad6a-adf5ef822c0c fwd="174.26.212.67,172.68.47.4" dyno=web.1 connect=0ms service=8ms status=304 bytes=93 protocol=https

2019-01-11T19:06:33.540637+00:00 heroku[router]: at=info method=GET path="/assets/application-7c9d81c598a0f0829fa038d6fd812dd9419dee8bee8f043b3ce91bede39ae1ac.js" host=www.crowdvest.co request_id=35a96ea1-8239-4c0b-8476-814c7b34d805 fwd="174.26.212.67,108.162.215.188" dyno=web.1 connect=0ms service=7ms status=304 bytes=93 protocol=https

2019-01-11T19:06:33.611516+00:00 heroku[router]: at=info method=GET path="/robots.txt" host=www.crowdvest.co request_id=eb3b039f-d7c3-42aa-8933-2592a1ab5c66 fwd="174.26.212.67,172.68.211.50" dyno=web.1 connect=1ms service=13ms status=200 bytes=421 protocol=https

2019-01-11T19:06:33.851308+00:00 heroku[router]: at=info method=GET path="/assets/application-5023ac9b44013d5b3cf2af3fe9e05d057eeecd06bb26840b7329d57c2b563d0b.css" host=www.crowdvest.co request_id=98574866-9a08-4ceb-a0fa-18533b4ce292 fwd="174.26.212.67,172.68.211.38" dyno=web.1 connect=0ms service=16ms status=304 bytes=93 protocol=https

2019-01-11T19:06:34.845904+00:00 heroku[router]: at=info method=GET path="/assets/team-5a12cd0e4da47b99ddb09c32617e93d74584977bec358c63ec41206f4f435b41.svg" host=www.crowdvest.co request_id=d14f6012-fe7b-46d4-ade6-338a037d77d8 fwd="174.26.212.67,172.68.46.111" dyno=web.1 connect=1ms service=4ms status=304 bytes=93 protocol=https

2019-01-11T19:06:34.846626+00:00 heroku[router]: at=info method=GET path="/assets/cash-e4d33605db031e67975684a412066b79573ea9a81e6189c6478abe784018f7ed.svg" host=www.crowdvest.co request_id=9c0ae46a-47b6-4cab-aef9-e50af2b28786 fwd="174.26.212.67,108.162.215.104" dyno=web.1 connect=1ms service=5ms status=304 bytes=93 protocol=https

2019-01-11T19:06:47.353299+00:00 heroku[router]: at=info method=GET path="/assets/matt-photo-cb54f0728df1f44b7ff7099eb8471b284a84df31738a24c7e44d1db9fffe1dfa.jpg" host=www.crowdvest.co request_id=11923866-e022-4276-b486-5fb513af778f fwd="174.26.212.67,173.245.48.163" dyno=web.1 connect=0ms service=4ms status=304 bytes=93 protocol=https

2019-01-11T19:06:47.285177+00:00 heroku[router]: at=info method=GET path="/assets/pie-chart-e7b01677e6d7da3f769c184fc5e7799b92127073405850e9301c191686ee09c7.svg" host=www.crowdvest.co request_id=ec74c40c-1165-4667-a079-9e9f6ae32fa3 fwd="174.26.212.67,162.158.58.157" dyno=web.1 connect=0ms service=2ms status=304 bytes=93 protocol=https

2019-01-11T19:06:47.353299+00:00 heroku[router]: at=info method=GET path="/assets/matt-photo-cb54f0728df1f44b7ff7099eb8471b284a84df31738a24c7e44d1db9fffe1dfa.jpg" host=www.crowdvest.co request_id=11923866-e022-4276-b486-5fb513af778f fwd="174.26.212.67,173.245.48.163" dyno=web.1 connect=0ms service=4ms status=304 bytes=93 protocol=https

2019-01-11T19:06:47.285177+00:00 heroku[router]: at=info method=GET path="/assets/pie-chart-e7b01677e6d7da3f769c184fc5e7799b92127073405850e9301c191686ee09c7.svg" host=www.crowdvest.co request_id=ec74c40c-1165-4667-a079-9e9f6ae32fa3 fwd="174.26.212.67,162.158.58.157" dyno=web.1 connect=0ms service=2ms status=304 bytes=93 protocol=https

2019-01-11T19:06:47.373824+00:00 heroku[router]: at=info method=GET path="/assets/hero-86eb20fc158ef4c774d0d2bae8587c6a284609e81aded79b299e49eb3a4ce913.jpg" host=www.crowdvest.co request_id=9966cd4d-e72b-4088-8aa1-2fb16cae7f56 fwd="174.26.212.67,173.245.48.139" dyno=web.1 connect=1ms service=6ms status=304 bytes=93 protocol=https

2019-01-11T19:06:47.357524+00:00 heroku[router]: at=info method=GET path="/assets/broker-check-7865b6db7a798d1b8e364e662cd6feb5d6098539260464394820395da3f7caad.png" host=www.crowdvest.co

Ready

☒ Autoscroll with Output

Shortcuts · Hide Logs



Summary

- ESTO Investments in SAR Processing are changing the way we interact with larger and larger volumes of data
- GPU-Accelerated generation of directly geocoded imagery enables previously prohibitive exploration of combinatoric interferometry
- GPU-Accelerated total workflow optimizations are enabling a degree of interactivity and experimentation on science problems
- A framework-based workflow visualizer, focused on products and interface simplicity, promises to make the interaction fun, iterative, and minable in the future